

Tracking Tensor Subspaces with Informative Random Sampling for Real-Time MR Imaging[†]

Morteza Mardani, Georgios B. Giannakis, and Kamil Ugurbil*

Submitted: September 15, 2016

Abstract

Magnetic resonance imaging (MRI) nowadays serves as an important modality for diagnostic and therapeutic guidance in clinics. However, the *slow acquisition* process, the dynamic deformation of organs, as well as the need for *real-time* reconstruction, pose major challenges toward obtaining artifact-free images. To cope with these challenges, the present paper advocates a novel subspace learning framework that permeates benefits from parallel factor (PARAFAC) decomposition of tensors (multiway data) to low-rank modeling of temporal sequence of images. Treating images as multiway data arrays, the novel method preserves spatial structures and unravels the latent correlations across various dimensions by means of the tensor subspace. Leveraging the spatio-temporal correlation of images, Tykhonov regularization is adopted as a rank surrogate for a least-squares optimization program. Alternating majorization minimization is adopted to develop online algorithms that recursively procure the reconstruction upon arrival of a new undersampled k -space frame. The developed algorithms are *provably convergent* and highly *parallelizable* with lightweight FFT tasks per iteration. To further accelerate the acquisition process, randomized subsampling policies are devised that leverage intermediate estimates of the tensor subspace, offered by the online scheme, to *randomly* acquire *informative* k -space samples. In a nutshell, the novel approach enables tracking motion dynamics under low acquisition rates ‘on the fly.’ GPU-based tests with real *in vivo* MRI datasets of cardiac cine images corroborate the merits of the novel approach relative to state-of-the-art alternatives.

Index Terms

[†] Work in this paper was supported by the NIH grant 1R01GM104975-01, and NSF grants 1514056 and 1500713. Parts of the paper was presented in the *36th IEEE Engineering in Medicine and Biology Conference (EMBC)*, Chicago, IL, September 2014; and the *23th ISMRM Annual Meeting and Exhibition*, Toronto, Ontario, May 2015.

* M. Mardani is with the Stanford University, School of Medicine, 875 Blake Wilbur Drive Stanford, CA 94305; Email: morteza@stanford.edu. G. B. Giannakis is with the Dept. of ECE and the Digital Technology Center, University of Minnesota, 200 Union Street SE, Minneapolis, MN 55455. Tel/fax:(612) 625-4287/625-2002; Email: georgios@umn.edu. K. Ugurbil is also with the Center for Magnetic Resonance Research and the University of Minnesota Medical School, 2021 Sixth street SE, Minneapolis, MN 55416, Tel:(612)626-9591; kamil@cmrr.umn.edu

Real-time MRI, PARAFAC decomposition, tensor subspace learning, low-rank, randomized subsampling.

I. INTRODUCTION

Since its inception in the 70s, magnetic resonance imaging (MRI) has emerged as a premier tool for biomedical imaging [15]. In recent years, MRI has also shown tremendous potential for dynamic processing. Through different protocols, dynamic MRI is able to provide images of tissues, perfusion, diffusion, spectroscopy, and susceptibility, both qualitatively as well as quantitatively in 3D high resolution, and in real time for every patient. The abundance of diverse data across time offers unprecedented opportunities to understand, diagnose, and treat diseases. Nevertheless, with such big blessings come big challenges, including: (c1) acquisition that is rather slow, allowing only for a limited amount of data to be collected per time slot to ensure that *motion* is frozen; and (c2) *real-time* image reconstruction is limited by the associated tradeoffs between speed and accuracy. Consequently, only low-resolution dynamic images can be acquired by state-of-the-art real-time MRI scanners.

Ample research has been carried out over the last decade to accelerate the dynamic MRI scanning process [9], [10], [12], [16], [17], [24], [25]. One way or another, existing works exploit the spatio-temporal correlation of MR images. Compressive sampling (CS) has been widely employed to leverage the parsimonious nature of data in a proper transform domain by means of sparsity and low rank regularization. The noteworthy representatives include k - t SPARSE [17], k - t GRAPPA [9], k - t SPARSE SENSE [24], k - t BLAST [12], k - t FOCUSS [10], k - t SLR [16], and k - t PCA [25]. They however rely on batch data processing, and typically non-smooth optimization modules which are relatively slow and demand high computational and storage resources for high resolution imaging. While batch processing is affordable for diagnostic purposes, image-guided therapeutic and surgical navigations demand real-time tracking for the organs of interest. There is a handful of studies on real-time MRI reconstruction that rely either on Kalman filtering, or, online compressive sampling; see e.g., [21], [28]. Kalman filtering based techniques [28] capture motion dynamics via state-space models, and end up with fast but low-quality reconstruction. CS-based methods such as [21] build on the motion sparsity and yield higher quality images but they are comparatively slow.

Aiming at fast and enhanced quality reconstruction, the present paper brings forth a novel tensor (multiway data array) subspace learning (TSL) framework that unravels the latent correlation structure of the MRI data stream. MR images comprise a multiway array with x, y, z coordinates, as well as time and possibly the coil dimension for parallel imaging. For a general linear observation of an M -way

array we postulate a low-rank model based on the parallel factor (PARAFAC) decomposition [11] of tensors that summarizes the tensor latent subspace in $M - 1$ factor matrices. A Tikhonov regularizer is adopted as a rank surrogate that regularizes a least-squares (LS) fitting cost to estimate the subspace and consequently interpolate the missing data. Broadening the scope of our precursor works in [23] and [22], and leveraging the decomposable structure of the sought optimization formulation, stochastic alternating minimization is adopted to develop iterative solvers for which the acquisition time coincides with the iteration index. Upon acquisition, the new datum with partial k -space data is first projected onto the latest subspace estimate, and the mismatch refines the subspace. The resulting procedure boils down to lightweight iterates with parallelized computations that suit GPU implementation for high-resolution imaging. In addition, the resulting subspace sequence is provably convergent to the stationary point set of the batch objective.

For the possibly parallel MRI, reconstruction schemes are introduced that either interpolate the misses in the k -space, or, directly retrieve in the image domain pursuing a tomographic approach. The proposed schemes offer real-time reconstruction of MR images ‘on the fly.’ Furthermore, the online subspace estimates, are utilized to devise a data-driven k -space subsampling rule to further accelerate the acquisition by collecting the most *informative* features for reconstruction. Specifically, inspired by randomized linear algebra approaches [18], we put forth a novel importance score that ranks the k -space entries to be acquired in the next frame, according to their coherence level with the latest tensor subspace. GPU-based simulated tests with two different *in vivo* cardiac cine MR image datasets corroborate the effectiveness of the novel reconstruction schemes in terms of speed, and motion tracking relative to k - t FOCUSS [10] and differential CS [21]. Last but not least, the scope of the proposed framework goes beyond dynamic MRI, and can indeed cater to other ‘big data’ inference tasks encountered with different medical imaging modalities.

The rest of this paper is organized as follows. Section II introduces preliminaries on tensor PARAFAC as well as rank regularization, and advocates a model to arrive at a generic optimization setup for reconstruction. Section III then develops iterative solvers to track the tensor subspace. Subsequently, Section IV focuses on dynamic and parallel MRI settings, where two reconstruction schemes are proposed to either interpolate misses in the k -space or image domain in a tomographic manner. Adaptive random subsampling to further accelerate MR acquisition process is the subject of Section V. Finally, real-data tests are reported in Section VI, while conclusions are drawn in Section VII.

Notation: Bold uppercase (lowercase) letters will denote matrices (column vectors with their entries in parenthesis), and calligraphic letters will be used for sets. Tensors or multi-way arrays are denoted by

bold, underlined uppercase letters. Operators $(\cdot)^\top$, $(\cdot)^*$, $(\cdot)^H$, $\text{tr}(\cdot)$, $\mathbb{E}[\cdot]$, $\sigma_{\max}(\cdot)$, \odot , and \circ will denote transposition, complex conjugate, Hermitian, matrix trace, statistical expectation, maximum singular value, Hadamard product, and outer product, respectively; $|\cdot|$ will be used for the cardinality of a set, and the magnitude of a scalar. The positive semidefinite matrix \mathbf{M} will be denoted by $\mathbf{M} \succeq \mathbf{0}$. The ℓ_p -norm of $\mathbf{x} \in \mathbb{R}^n$ is $\|\mathbf{x}\|_p := (\sum_{i=1}^n |x_i|^p)^{1/p}$ for $p \geq 1$. For two matrices $\mathbf{M}, \mathbf{U} \in \mathbb{R}^{n \times p}$, $\langle \mathbf{M}, \mathbf{U} \rangle := \text{tr}(\mathbf{M}^\top \mathbf{U})$ denotes their trace inner product, and $\|\mathbf{M}\|_F := \sqrt{\text{tr}(\mathbf{M}\mathbf{M}^\top)}$ is the Frobenius norm. The $n \times n$ identity matrix will be represented by \mathbf{I}_n , while $\mathbf{0}_n$ will stand for the $n \times 1$ vector of all zeros, $\mathbf{0}_{n \times p} := \mathbf{0}_n \mathbf{0}_p^\top$, and $[n] := \{1, 2, \dots, n\}$. Also, \mathbf{e}_i denotes the canonical vector with one at i -th entry and zero elsewhere, while the operators vec and unvec stack the columns of a matrix on top of each other, and vice versa.

II. PRELIMINARIES AND PROBLEM STATEMENT

As modern and massive datasets become increasingly complex and heterogeneous, in many application setups one encounters data structures indexed by three or more variables giving rise to a tensor, instead of just two variables indexing data organized in a matrix. A few examples of time-indexed, medical tensor data include [1]: (i) images acquired in parallel MRI across various coils, as well as snapshots across time and patients, collected in a five-dimensional array with (phase encoding, frequency encoding, coil, time, patient); and (ii) Electroencephalograms (EEGs), where the signal of each electrode is a time-frequency matrix; thus, data from multiple channels is three-dimensional (temporal, spectral, and spatial) and may be incomplete if electrodes become loose or disconnected for a period of time.

A. Low-rank PARAFAC decomposition

For multiple, say $M \geq 2$, vectors $\mathbf{a}_m \in \mathbb{C}^{N_m \times 1}$, the outer product $\mathbf{a}_1 \circ \dots \circ \mathbf{a}_M$ is an $N_1 \times \dots \times N_M$ rank-one M -way array with (n_1, \dots, n_M) -th entry given by $\prod_{m=1}^M a_{n_m, m}$, where $a_{n_m, m}$ is the n_m -th entry of \mathbf{a}_m . This generalizes the matrix case ($M = 2$), where $\mathbf{a}_1 \circ \mathbf{a}_2 = \mathbf{a}_1 \mathbf{a}_2^\top$ is a rank-one matrix. The rank of a tensor $\underline{\mathbf{X}}$ is defined as the minimum number of outer products required to synthesize $\underline{\mathbf{X}}$. The PARAFAC model is arguably the most basic model because of its direct relationship to tensor rank. Specifically, it is natural to form an R -rank approximation of tensor $\underline{\mathbf{X}} \in \mathbb{C}^{N_1 \times \dots \times N_M}$ as

$$\underline{\mathbf{X}} \approx \sum_{r=1}^R \mathbf{a}_r^{(1)} \circ \dots \circ \mathbf{a}_r^{(M)}. \quad (1)$$

When the approximation is exact, (1) is the PARAFAC decomposition of $\underline{\mathbf{X}}$. Accordingly, the minimum value R for which the exact decomposition is possible is (by definition) the rank of $\underline{\mathbf{X}}$. Different from




Fig. 1. A rank- R PARAFAC decomposition of the three-way tensor $\underline{\mathbf{X}}$.

the matrix case, there is no straightforward algorithm to determine the rank of a given tensor, a problem that is known to be NP-hard [11].

With reference to (1), introduce the factor matrices $\mathbf{A}_m := [\mathbf{a}_1^{(m)}, \dots, \mathbf{a}_R^{(m)}] \in \mathbb{C}^{N_m \times R}$, $i \in [M]$, and let $\underline{\mathbf{X}}_\ell^{(k)}$, $\ell \in [N_k]$ denote the ℓ -th slice of $\underline{\mathbf{X}}$ along its k -th mode, such that $\underline{\mathbf{X}}_\ell^{(k)}(n_1, \dots, n_{k-1}, n_{k+1}, \dots, n_M) = \underline{\mathbf{X}}(n_1, \dots, \ell, \dots, n_M)$. This ℓ -th slice can then be expressed as (cf. (1))

$$\underline{\mathbf{X}}_\ell^{(k)} = \sum_{r=1}^R \gamma_{\ell,r}^{(k)} \mathbf{a}_r^{(1)} \circ \dots \circ \mathbf{a}_r^{(k-1)} \circ \mathbf{a}_r^{(k+1)} \circ \dots \circ \mathbf{a}_r^{(M)}, \quad \ell = 1, \dots, N_k \quad (2)$$

where $\gamma_\ell \in \mathbb{C}^R$ denotes the ℓ -th row of \mathbf{A}_k whose r -th column is $\mathbf{a}_r^{(k)}$. To gain intuition, imagine a three-way data array ($M = 3$), where the slices form matrices, and thus the ℓ -th slice across the tube dimension is given by $\mathbf{X}_\ell = \sum_{r=1}^R \gamma_{\ell,r}^{(3)} \mathbf{a}_r^{(1)} \mathbf{a}_r^{(2)\top}$. It is apparent that a slice \mathbf{X}_ℓ can be represented as a linear combination of R rank-one matrices $\{\mathbf{a}_r^{(1)} \mathbf{a}_r^{(2)\top}\}_{r=1}^R$, which constitute the bases for the tensor tube subspace. In the same manner, one can argue that for a general M -th order tensor, the rank-one tensors $\{\mathbf{a}_r^{(1)} \circ \dots \circ \mathbf{a}_r^{(k-1)} \circ \mathbf{a}_r^{(k+1)} \circ \dots \circ \mathbf{a}_r^{(M)}\}_{r=1}^R$ form the bases for the tensor subspace along the k -th mode. Likewise, R -dimensional vector $\gamma_\ell^{(k)}$ collects the tensor principal components.

This study aims at discovering the latent subspace that is captured by the matrices $\{\mathbf{A}_m\}_{m=1, m \neq k}^M$. This will become handy later. Given $\underline{\mathbf{X}}$, under mild conditions, matrices $\{\mathbf{A}_m\}_{m=1}^M$ are unique up to a common column permutation and scaling (meaning that PARAFAC is essentially identifiable for $M \geq 3$); see e.g. [4], [13]. It is worth commenting that the factor matrices $\{\mathbf{A}_m\}_{m=1}^M$ are not necessarily orthogonal, and may even be rank deficient. Thanks to its essential uniqueness, PARAFAC has become the model of choice when one is primarily interested in revealing latent structure in multiway data arrays. Considering the analysis of a dynamic social network for instance, each of the rank-one factors could correspond to communities that e.g., persist or form and dissolve dynamically across time [2].

PARAFAC's link with the tensor rank can be used to postulate low-rank tensor models. However, as mentioned earlier even finding the tensor rank is an NP-hard problem. Parallel to the matrix nuclear-norm, tractable surrogates can be adopted for the tensor PARAFAC-rank that approximates the rank through

the norm of its factor matrices. One such surrogate for $M = 3$ is introduced in our companion work [3]. Generalizing [3] to multi-way arrays, we adopt the following rank surrogate for the tensor rank of $\underline{\mathbf{X}}$

$$\mathcal{Q}(\underline{\mathbf{X}}) := \min_{\{\mathbf{A}_i \in \mathbb{C}^{N_i \times R}\}_{i=1}^M} \frac{1}{M} \sum_{i=1}^M \|\mathbf{A}_i\|_F^2 \quad \text{s.t.} \quad \underline{\mathbf{X}} = \sum_{r=1}^R \mathbf{a}_r^{(1)} \circ \dots \circ \mathbf{a}_r^{(M)}. \quad (3)$$

The proof that $\mathcal{Q}(\underline{\mathbf{X}})$ induces low tensor rank follows the main ideas in [3] for $M = 3$, but it is omitted here due to space limitation. Nonetheless, the formal claim is given next.

Lemma 1: *If $\sigma_r := \Pi_{m=1}^M \|\mathbf{a}_r^{(m)}\|$ denotes the r -th singular value of the M -way tensor $\underline{\mathbf{X}}$, it then holds that*

$$\mathcal{Q}(\underline{\mathbf{X}}) = \left(\sum_{r=1}^R |\sigma_r|^{2/M} \right)^{2/M}.$$

Lemma 1 asserts that $\mathcal{Q}(\underline{\mathbf{X}})$ promotes sparsity across the singular values of $\underline{\mathbf{X}}$. Note that for $M = 2$ the adopted regularizer boils down to the well-known matrix nuclear norm [7].

B. Subsampled data model

In various application domains, the physical data of interest collected in a tensor vary slowly and can thus be approximated by a stationary process. Accordingly, consider that the M -way tensor process $\{\underline{\mathbf{L}}_t\}$ lives in a low-dimensional subspace \mathcal{L} . From this process, undersampled observations $\{\mathbf{y}_t \in \mathbb{C}^{L_t}\}$ that are streaming over time obey the model

$$y_t^{(\ell)} = \langle \underline{\mathbf{L}}_t, \underline{\mathbf{W}}_t^{(\ell)} \rangle + v_t^{(\ell)}, \quad \ell = 1, \dots, L_t \quad (4)$$

where the projection tensor $\underline{\mathbf{W}}_t^{(\ell)} \in \mathbb{C}^{N_1 \times \dots \times N_{M-1}}$ sketches (meaning subsamples) $\underline{\mathbf{L}}_t$, and $v_t^{(\ell)}$ accounts for the errors and unmodeled dynamics. This model hits several modern application domains such as dynamic MRI, where the ground-truth sequence of images form a three-way data cube, and per time slot t a small subset of k -space data in the Fourier domain are acquired; further details are provided in Section IV.

Assume temporarily that one has only access to a batch of observations (4) during the time horizon $t \in [1, T]$. Collect the $(M - 1)$ -way tensors $\{\underline{\mathbf{L}}_t\}_{t=1}^T$ into a larger M -way tensor $\underline{\mathbf{L}}$ with the M -th mode representing time. Low dimensionality implies $\underline{\mathbf{L}}$ has low PARAFAC rank. All in all, we wish to identify $\{\underline{\mathbf{L}}_t\}_{t=1}^T$ given the observations $\{\mathbf{y}_t\}_{t=1}^T$ along with the projection tensors $\{\underline{\mathbf{W}}_t^{(\ell)}\}_{t=1}^T$, assuming $\underline{\mathbf{L}}$ is a

low-rank tensor. Letting $\gamma_t \in \mathbb{C}^R$ denote the t -th row of $\mathbf{A}_M[t]$, a natural estimator $\hat{\underline{\mathbf{L}}}$ is

$$(P1) \quad \hat{\underline{\mathbf{L}}} = \arg \min_{\{\mathbf{A}_m\}_{m=1}^M, \{\gamma_t\}_{t=1}^T} \frac{1}{2} \sum_{t=1}^T \sum_{\ell=1}^{L_t} (y_t^{(\ell)} - \langle \underline{\mathbf{L}}_t, \underline{\mathbf{W}}_t^{(\ell)} \rangle)^2 + \frac{\lambda}{2} \sum_{m=1}^M \|\mathbf{A}_m\|_F^2$$

$$\text{s. to} \quad \underline{\mathbf{L}}_t = \sum_{r=1}^R \gamma_{t,r} \mathbf{a}_r^{(1)} \circ \dots \circ \mathbf{a}_r^{(M-1)}, \quad t = 1, \dots, T$$

which fits the data to the postulated model (4) in the LS sense, and promotes low tensor rank through the regularizer $\sum_{m=1}^M \|\mathbf{A}_m\|_F^2$. Tuning λ , controls the desired rank level. Note all the data and optimization variables in (P1) are complex valued.

In big data settings, the ambient dimensions $\{N_m\}_{m=1}^{M-1}$ can be quite large, and the tensor slices can be streaming over time with possibly $N_M \rightarrow \infty$ streams collected over time. Before delving into online solvers of (P1) for streaming observations, a couple of noteworthy properties of (P1) are in order. First, the rank regularization avoids the scaling ambiguity associated with the multilinear terms as stated next.

Lemma 2: *Every stationary point of (P1) returns a tensor subspace with equal norm bases; that is, $\|\mathbf{a}_r^{(m)}\| = \|\mathbf{a}_r^{(m')}\|, \forall m, m' \in [M-1]$, and $r \in [R]$.*

Proof: It readily follows by equating the gradient of (P1)'s objective w.r.t. $\mathbf{a}_r^{(m)}$ (see also (8)) to zero, and taking the inner product of both sides with $\mathbf{a}_r^{(m)}$. ■

Equal norm bases fix the scaling ambiguity inherent to the PARAFAC model. It also implies that the tensor singular values in Lemma 1 are simply expressed as $\sigma_r = \|\mathbf{a}_r^{(1)}\|^M$.

For large-scale inference tasks with $\prod_{m=1}^{M-1} N_m$ large solving (P1) incurs prohibitive complexity and storage to implement in batch mode. In addition, certain applications demand real-time processing upon acquisition of a new datum based on the past and current data, namely $\{y_\tau^{(\ell)}, \ell \in [L_t]\}_{\tau=1}^t$. In essence, (P1) involves $R(N_1 + \dots + N_{M-1} + t)$ variables associated with the low-rank components, which can grow prohibitively with t , and eventually exceed the storage and computational limits. These obstacles press the need for online iterative solvers that can acquire data sequentially and perform simple update tasks. The ensuing section introduces machinery to arrive at such efficient online solvers.

III. TRACKING TENSOR SUBSPACE

As elaborated in Section II-B the low-rank tensor $\underline{\mathbf{L}}_t$ lies in a low-dimensional subspace $\mathcal{L} \subset \mathbb{C}^{N_1 \times \dots \times N_{M-1}}$. With reference to PARAFAC-rank, \mathcal{L} is characterized by a small number R of rank-one $(M-1)$ -way arrays $\{\mathbf{a}_r^{(1)} \circ \dots \circ \mathbf{a}_r^{(M-1)}\}_{r=1}^R$, captured by the factor matrices $\{\mathbf{A}_m\}_{m=1}^{M-1}$. Learning these time-invariant factor matrices is the first step towards reconstructing the low-rank tensor of interest. With the streaming

observations, at t -th acquisition time one is given $T = t$ data snapshots, and accordingly (with a slight abuse of notation) one is motivated to recast (P1) in the separable form

$$(P2) \quad \min_{\{\mathbf{A}_m\}_{m=1}^{M-1}} \frac{1}{2t} \sum_{\tau=1}^t \min_{\gamma_\tau} \left\{ \sum_{\ell=1}^{L_t} \left(y_\tau^{(\ell)} - \sum_{r=1}^R \gamma_{\tau,r} \langle \mathbf{a}_r^{(1)} \circ \dots \circ \mathbf{a}_r^{(M-1)}, \mathbf{W}_t^{(\ell)} \rangle \right)^2 + \frac{\lambda}{2} \|\gamma_\tau\|^2 \right\} + \frac{\lambda}{2t} \sum_{m=1}^{M-1} \|\mathbf{A}_m\|_F^2.$$

Apparently, finding the optimal solution of the nonconvex program (P2) becomes computationally challenging especially for M large. Hence, approximations that can afford simple iterative updates while approaching the optimal solution are well motivated. One such approximation for online rank minimization leveraging the separable nuclear-norm regularization (3) was introduced in [22] for matrices ($M = 2$), in the context of unveiling network anomalies, and in [23] for imputation of three-way tensors. Building on [23] and [22], online solvers are developed next for the general $M \geq 3$ case.

A. Stochastic alternating minimization

Towards deriving a real-time, computationally efficient, and recursive solver of (P2), an alternating-minimization (AM) method is adopted in which iterations coincide with the index t of data acquisition. In accordance with (P2), consider the instantaneous regularized LS cost

$$f_t(\{\mathbf{A}_i\}_{i=1}^{M-1}; \gamma_t) := \frac{1}{2} \sum_{\ell=1}^{L_t} \left(y_t^{(\ell)} - \sum_{r=1}^R \gamma_{t,r} \langle \mathbf{a}_r^{(1)} \circ \dots \circ \mathbf{a}_r^{(M-1)}, \mathbf{W}_t^{(\ell)} \rangle \right)^2 + (\lambda/2t) \sum_{m=1}^{M-1} \|\mathbf{A}_m\|_F^2. \quad (5)$$

The iterative procedure adopted here consists of two major steps. The first step (S1) relies on the recently updated subspace, namely $\{\mathbf{A}_m[t-1]\}_{m=1}^{M-1}$, to solve the inner optimization, which yields the principal components $\gamma_t = \arg \min_{\gamma} f(\{\mathbf{A}_i\}_{i=1}^{M-1}; \gamma)$. In the second step (S2), the tensor subspace \mathcal{L}_t is updated by moving $\{\mathbf{A}_m\}_{m=1}^{M-1}$ along the opposite direction of the gradient, namely $-\nabla f_t(\{\mathbf{A}_i\}_{i=1}^{M-1}; \gamma_t)$.

For (S1), collect $y_t^{(\ell)}$ in $\mathbf{y}_t \in \mathbb{C}^{L_t}$, and define matrix $\Phi_t := [\phi_1^{(t)}, \dots, \phi_{L_t}^{(t)}]^\top \in \mathbb{C}^{L_t \times R}$, where $[\phi_t^{(\ell)}]_r := \langle \mathbf{a}_r^{(1)} \circ \dots \circ \mathbf{a}_r^{(M-1)}, \mathbf{W}_t^{(\ell)} \rangle$. The projection of \mathbf{y}_t onto the low-dimensional subspace \mathcal{L}_t is then obtained by solving the LS ridge-regression problem

$$\gamma_t = \arg \min_{\gamma \in \mathbb{C}^R} \frac{1}{2} \|\mathbf{y}_t - \Phi_t \gamma\|^2 + \frac{\lambda}{2} \|\gamma\|^2$$

which admits the closed-form solution $\gamma_t = (\Phi_t^\top \Phi_t + \lambda \mathbf{I}_R)^{-1} \Phi_t^\top \mathbf{y}_t$ that depends linearly on the subsampled data. To avoid the $R \times R$ matrix inversion, consider the SVD $\Phi_t = \mathbf{U}_t \Sigma_t \mathbf{V}_t$ to end up with $\gamma_t = \mathbf{V}_t \Sigma_t^{-1} \mathbf{D}_t \mathbf{U}_t^\top \mathbf{y}_t$, where $\mathbf{D}_t \in \mathbb{C}^{R \times R}$ is a diagonal matrix with $[\mathbf{D}_t]_{i,i} = \sigma_i^2 / (\sigma_i^2 + \lambda_*)$.

The second step (S2) deals with updating the factor matrices given $\{\gamma_\tau\}_{\tau=1}^t$ by solving

$$\{\mathbf{A}_m[t]\}_{m=1}^{M-1} = \arg \min_{\{\mathbf{A}_m\}_{m=1}^{M-1}} C_t(\{\mathbf{A}_m\}_{m=1}^{M-1}) := \frac{1}{t} \sum_{\tau=1}^t f_\tau(\{\mathbf{A}_m\}_{m=1}^{M-1}; \gamma_\tau). \quad (6)$$

Apparently, (6) is a nonconvex program for $M \geq 3$ due to the multilinear terms in the LS cost, and is thus tough to solve optimally. To mitigate this computational challenge, consider the following quadratic approximant of f_t

$$\begin{aligned} \tilde{f}_t(\{\mathbf{A}_m\}_{m=1}^{M-1}; \gamma_t) &= f_t(\{\mathbf{A}_m[t-1]\}_{m=1}^{M-1}; \gamma_t) \\ &+ \sum_{m=1}^{M-1} \langle \nabla_{\mathbf{A}_m} f_t(\{\mathbf{A}_m[t-1]\}_{m=1}^{M-1}; \gamma_t), \mathbf{A}_m - \mathbf{A}_m[t-1] \rangle + \frac{\alpha_t}{2} \sum_{m=1}^{M-1} \|\mathbf{A}_m - \mathbf{A}_m[t-1]\|_F^2 \end{aligned}$$

where $\alpha_t \geq \max_m \left\{ \sigma_{\max} [\nabla_{\mathbf{A}_m}^2 f_t(\{\mathbf{A}_m[t-1]\}_{m=1}^{M-1}; \gamma_t)] \right\}$. With regards to the surrogate \tilde{f}_t , it is useful to recognize that it is locally tight, meaning that (i) $f_t(\{\mathbf{A}_m[t-1]\}_{m=1}^{M-1}; \gamma_t) = \tilde{f}_t(\{\mathbf{A}_m[t-1]\}_{m=1}^{M-1}; \gamma_t)$, and similarly $\nabla f_t(\{\mathbf{A}_m[t-1]\}_{m=1}^{M-1}; \gamma_t) = \nabla \tilde{f}_t(\{\mathbf{A}_m[t-1]\}_{m=1}^{M-1}; \gamma_t)$; and (ii) it upper bounds f_t , that is $f_t(\{\mathbf{A}_m\}_{m=1}^{M-1}; \gamma_t) \leq \tilde{f}_t(\{\mathbf{A}_m\}_{m=1}^{M-1}; \gamma_t)$, for all $\mathbf{A}_m \in \mathbb{C}^{N_m \times R}$, and $m \in [M-1]$.

Apart from tightness, separability across factors is another attractive feature of \tilde{f}_t because it allows for parallel implementation. Plugging in \tilde{f}_t into the cost C_t yields $\tilde{C}_t := (1/t) \sum_{\tau=1}^t \tilde{f}_\tau$, the minimizer of which is obtained (after equating the gradient to zero) as

$$\mathbf{A}_m[t] = \frac{1}{\bar{\alpha}_t} \sum_{\tau=1}^t \alpha_\tau \left\{ \mathbf{A}_m[\tau-1] - \alpha_\tau \nabla_{\mathbf{A}_m} f_\tau(\{\mathbf{A}_m[\tau-1]\}_{m=1}^{M-1}; \gamma_\tau) \right\}$$

where $\bar{\alpha}_t := \sum_{\tau=1}^t \alpha_\tau$. After rearranging terms one arrives at the recursion

$$\begin{aligned} \mathbf{A}_m[t] &= \underbrace{\left(\frac{1}{\bar{\alpha}_t} \right) \sum_{\tau=1}^{t-1} \alpha_\tau \left\{ \mathbf{A}_m[\tau-1] - \alpha_\tau^{-1} \nabla_{\mathbf{A}_m} f_\tau(\{\mathbf{A}_m[\tau-1]\}_{m=1}^{M-1}; \gamma_\tau) \right\}}_{:= \bar{\alpha}_{t-1} \mathbf{A}_m[t-1]} \\ &+ \left(\frac{\alpha_t}{\bar{\alpha}_t} \right) \left\{ \mathbf{A}_m[t-1] - \alpha_t \nabla_{\mathbf{A}_m} f_t(\{\mathbf{A}_m[t-1]\}_{m=1}^{M-1}; \gamma_t) \right\} \\ &= \mathbf{A}_m[t-1] - (\bar{\alpha}_t)^{-1} \nabla_{\mathbf{A}_m} f_t(\{\mathbf{A}_m[t-1]\}_{m=1}^{M-1}; \gamma_t), \quad m \in [M-1]. \end{aligned} \quad (7)$$

Interestingly, (7) is nothing but a single stochastic gradient descent step.

The gradient is separable across columns of \mathbf{A}_m . Considering it w.r.t. each basis vector $\mathbf{a}_r^{(m)}$ leads to the closed-form expression

$$\begin{aligned} \nabla_{\mathbf{a}_r^{(m)}} f_t(\mathbf{A}_1, \dots, \mathbf{A}_{M-1}) &= (\lambda/t) \mathbf{a}_r^{(m)} \\ &- \sum_{\ell=1}^L \gamma_{t,r}^* \left(y_t^{(\ell)} - \sum_{r=1}^R \gamma_{t,r} \langle \mathbf{a}_r^{(1)} \circ \dots \circ \mathbf{a}_r^{(M-1)}, \mathbf{w}_t^{(\ell)} \rangle \right) \left(\mathbf{w}_t^{(\ell)} \times_{i=1, i \neq m}^{M-1} \mathbf{a}_r^{(i)} \right)^* \end{aligned} \quad (8)$$

where \times_i denotes the Tucker mode- i product [11]. It is also useful to recognize that the Hessian of f_t admits the simple form

$$\begin{aligned} \nabla_{\mathbf{a}_r}^2 f_t(\mathbf{A}_1, \dots, \mathbf{A}_{M-1}) &= (\lambda/t) \mathbf{I}_{N_1} \\ &+ |\gamma_{t,r}|^2 \sum_{\ell=1}^{L_t} \left(\underline{\mathbf{W}}_t^{(\ell)} \times_{i=1, i \neq m}^{M-1} \mathbf{a}_r^{(i)} \right) \left(\underline{\mathbf{W}}_t^{(\ell)} \times_{i=1, i \neq m}^{M-1} \mathbf{a}_r^{(i)} \right)^H. \end{aligned} \quad (9)$$

All in all, the gradient iterations for learning the tensor subspace proceed in parallel as follows

$$\mathbf{a}_r^{(m)}[t] = \mathbf{a}_r^{(m)}[t-1] - \mu_t \nabla_{\mathbf{a}_r^{(m)}} f_t(\{\mathbf{A}_m[t-1]\}_{m=1}^{M-1}; \gamma_t), \quad r \in [R], \quad m \in [M-1] \quad (10)$$

where $\mu_t = \bar{\alpha}_t^{-1}$ denotes the step size. The resulting algorithm is listed under Table 1.

B. Convergence analysis

Convergence of the first-order subspace iterates in Algorithm 1 is granted following the analysis developed in [19] for convergence of dictionary learning, and our precursors in [22], [23] establishing subspace convergence for imputation of two- and three-way arrays. The proof relies on martingale sequences, and in order to render the analysis tractable, it adopts the following assumptions:

(as1) The data stream $\{\mathbf{y}_t\}$ forms an i.i.d. random process that is uniformly bounded; and

(as2) The subspace updates $\{\mathcal{L}[t]\}$ where $\mathcal{L}[t] := \{\mathbf{A}_m[t]\}_{m=1}^{M-1}$ lies in a compact set.

With these assumptions, the convergence claim is formalized as follows.

Proposition 1: *If the subspace iterates $\{\mathcal{L}[t]\}$ lie in a compact set, and the step-size sequence $\{\mu_t\}$ satisfies $(\mu_t)^{-1} := \sum_{\tau=1}^t \alpha_\tau \geq ct$, $\forall t$ for some $c > 0$, where $c' \geq \alpha_t \geq \max_m \left\{ \sigma_{\max} [\nabla_{\mathbf{A}_m}^2 f_t(\{\mathbf{A}_m[t-1]\}_{m=1}^{M-1}; \gamma_t)] \right\}$, $\forall t$ for some $c' > 0$, then $\lim_{t \rightarrow \infty} \nabla C_t(\mathcal{L}[t]) \rightarrow \mathbf{0}$, a.s.; i.e., the tensor subspace iterates $\{\mathcal{L}[t]\}$ asymptotically coincide with the stationary point set of the batch program (P2).*

Note the step-size controlling constants c and c' determine the speed of convergence, and are chosen according to the observation parameters including the projection tensors $\{\underline{\mathbf{W}}_t^{(\ell)}\}$. In particular, c is related to the degree of curvature of the instantaneous loss f_t , and in a similar manner c' is tied to the smoothness of f_t , and admits small values when the acquired observations end up with a smooth loss function having a small Lipschitz constant.

Remark 1 [Parallelizable updates]: Implementing Algorithm 1 per time instant t involves updating the projection coefficients as well as the tensor subspace. The former mainly entails SVD computation of a $L_t \times R$ size matrix Φ_t , which incurs $\mathcal{O}(L_t R^2)$ operations. The latter is also nicely parallelizable across both the basis index r and the factor index m , and hence MR updates can be carried out simultaneously via parallel processors such as GPUs.

Algorithm 1 Online rank-regularized tensor subspace learning

input $\{y_t^{(\ell)}, \mathbf{W}_t^{(\ell)}, \ell \in [L_t]\}_{t=1}^\infty, \{\mu_t\}_{t=1}^\infty, \lambda, R, M$.

initialize $\{\mathbf{A}_m[1]\}_{m=1}^M$ at random.

for $t = 1, 2, \dots$ **do**

(S1) Projection coefficients update

$[\Phi_t]_{\ell,r} = \langle \mathbf{a}_r^{(1)}[t-1] \circ \dots \circ \mathbf{a}_r^{(M-1)}[t-1], \mathbf{W}_t^{(\ell)} \rangle, \Phi_t = \mathbf{U}_t \Sigma_t \mathbf{V}_t^\top$

$\mathbf{D}_t = \text{diag}[\sigma_1(\sigma_1^2 + \lambda)^{-1}, \dots, \sigma_R(\sigma_R^2 + \lambda)^{-1}]$

$\gamma_t = \mathbf{V}_t \Sigma_t^{-1} \mathbf{D}_t \mathbf{U}_t^\top \mathbf{y}_t$

$e_t^{(\ell)} := y_t^{(\ell)} - \langle \phi_t^{(\ell)}, \gamma_t \rangle$

(S2) Parallel subspace update $[(m, r) \in [M] \times [R]]$

$\mathbf{a}_r^{(m)}[t] = (1 - \mu_t \lambda / t) \mathbf{a}_r^{(m)}[t-1] + \mu_t \gamma_{t,r}^* \left(\sum_{\ell=1}^{L_t} e_t^{(\ell)} (\mathbf{W}_t^{(\ell)})^* \right) \times_{i=1, i \neq m}^{M-1} (\mathbf{a}_r^{(i)}[t-1])^*$

return $(\{\mathbf{A}_m[t]\}_{m=1}^{M-1}, \gamma_t)$

end for

The ensuing section deals with application of the proposed tensor subspace learning scheme for reconstructing dynamic images in time-resolved MRI.

IV. ACCELERATING DYNAMIC MRI

Dynamic MRI acquires a low-spatial yet high-temporal resolution sequence of images, which renders a possibly sizable portion of measurements per snapshot inaccurate or missing. Fortunately, but the temporal correlation of images can be leveraged to interpolate these misses.

MRI typically uses a phased array of coils, each imaging a limited spatial region of the object. To begin, consider that single-coil MRI is used to acquire the ground-truth image sequence $\{\mathbf{L}_t\}_{t=1}^T$ with the complex-valued image $\mathbf{L}_t \in \mathbb{C}^{N_1 \times N_2}$ corresponding to t -th frame. Entries of \mathbf{L}_t record the collective magnetic field (including both phase and magnitude) induced per tissue voxel. Data corresponding to the t -th frame acquired in the frequency domain (hereafter referred to as k -space data) are then modeled as

$$y_t^{(i,j)} = [\mathcal{F}(\mathbf{L}_t)]_{i,j} + v_{i,j}, \quad (i, j) \in \Omega_t \quad (11)$$

where $\mathcal{F}(\cdot)$ denotes the two-dimensional discrete Fourier transform (DFT) operator, and the set $\Omega_t \subset [N_1] \times [N_2]$ indexes the k -space data support. The acquisition time is clearly proportional to the sample count $\sum_{\tau=1}^t |\Omega_\tau|$, and it is desired to be as small as possible.

In what follows, a tensor imputation based approach is introduced first based on k -space correlations to interpolate the misses. For the multicoil scenario then a tomographic reconstruction scheme is devised

that exploits the low rank directly on the image domain. To set up notation, the two-dimensional DFT operator $\mathcal{F}(\cdot)$ can be written in compact matrix form as $\mathcal{F}(\mathbf{X}_t) = \mathbf{F}_l \mathbf{X}_t \mathbf{F}_r$, with the left and right matrices $\mathbf{F}_l \in \mathbb{R}^{N_1 \times N_1}$ and $\mathbf{F}_r \in \mathbb{R}^{N_2 \times N_2}$, respectively, denoting orthonormal symmetric DFT matrices. The subsampling matrix in (4) is then $\mathbf{W}_t^{(\ell)} = \mathbf{F}_l \mathbf{e}_i \mathbf{e}_j^\top \mathbf{F}_r$, and hence the projection corresponding to the (i, j) -th DFT coefficient is $\langle \mathbf{X}_t, \mathbf{W}_t^{(\ell)} \rangle = [\mathcal{F}(\mathbf{X}_t)]_{i,j}$.

A. k -space Interpolation

Recall that $\{\mathbf{L}_t\}$ is the underlying image sequence of interest that lies in a low-rank tensor subspace \mathcal{L} . Being an *orthonormal transform*, the Fourier operator preserves dimensionality, meaning that $\{\mathbf{X}_t = \mathcal{F}(\mathbf{L}_t)\}$ lies in a linear tensor subspace, say $\mathcal{L}_F \subset \mathbb{C}^{N_1 \times N_2}$ with $\dim(\mathcal{L}_F) = \dim(\mathcal{L})$. Note also that the magnitude of eigenvalues for both the image and k -space data remains the same.

Accordingly, one can build on the low PARAFAC rank of the three-way tensor $\underline{\mathbf{X}}$ to interpolate the misses from the present k -space entries. In this direction, given the partial k -space measurements

$$y_t^{(i,j)} = x_t^{(i,j)} + v_t^{(i,j)}, \quad (i, j) \in \Omega_t$$

we postulate a trilinear model $x_t^{(i,j)} = \langle \alpha_i, \beta_j, \gamma_t \rangle$ with $\alpha_i, \beta_j, \gamma_t$ denoting rows of $\mathbf{A}_1, \mathbf{A}_2, \mathbf{A}_3$, respectively. Choosing the sketching operator $\mathbf{W}_t^{(i,j)} = \mathbf{e}_i \mathbf{e}_j^\top$, one can solve (P2) to learn the tensor factor matrices $\mathbf{A}_1[t]$ and $\mathbf{A}_2[t]$ ‘on the fly,’ and subsequently interpolate to obtain $\hat{\mathbf{X}}_t := \mathbf{A}_1[t] \text{diag}(\gamma_t) \mathbf{A}_2^\top[t]$. With the imputed matrix $\hat{\mathbf{X}}_t$ at hand, the ground-truth image can then be reconstructed using the magnitude of $\mathcal{F}^{-1}(\hat{\mathbf{X}}_t)$, namely $[\hat{\mathbf{L}}_t]_{i,j} = |[\mathcal{F}^{-1}(\hat{\mathbf{X}}_t)]_{i,j}|$.

The corresponding algorithm specialized to the MRI task is listed under Algorithm 2. It can be seen as a special case of the general Algorithm 1 upon fixing $[\Phi_t]_{\ell,r} = [\mathbf{a}_r^{(1)}[t]]_i [\mathbf{a}_r^{(2)}[t]]_j$ for $\ell \equiv (i, j)$. The iterations admit a simple and interpretable form, where the i and j rows of $\mathbf{A}_1[t]$ and $\mathbf{A}_2[t]$, respectively, are updated once the k -space datum $(i, j) \in \Omega_t$ arrives.

This interpolation-based approach is particularly attractive when one can split each $N_1 \times N_2$ k -space image into $K_1 \times K_2$ (non)overlapping patches of size $n_1 \times n_2$, with $K_1 = N_1/n_1$ and $K_2 = N_2/n_2$. Patching is known to better leverage the local image features [27]. Patches must be sufficiently sizable to preserve the spatio-temporal correlations, and form a tensor with low rank ρ . This idea reduces the variable count associated with the subspace from $(N_1 + N_2)R$ to $(n_1 + n_2)\rho$ which can lead to significant computational savings. In addition, the large number of frames facilitates learning the tensor subspace, especially for MRI scans with low temporal resolution. Moreover, subsampling strategies, discussed in Section V, are immediately applicable to reduce acquisition time. In the multi-coil acquisition scenario

Algorithm 2 Online interpolation-based MRI

input $\{y_t^{(i,j)}, \ell \in \Omega_t\}_{t=1}^\infty, \{\mu_t\}_{t=1}^\infty, \lambda, R$.
initialize $(\mathbf{A}_1[0], \mathbf{A}_2[0])$.
for $t = 1, 2, \dots$ **do**
 (S1) Projection coefficients update
 $[\Phi_t]_{\ell,r} = [\mathbf{a}_r^{(1)}[t-1]]_i [\mathbf{a}_r^{(2)}[t-1]]_j, \ell \equiv (i,j) \in \Omega_t, \quad \Phi_t = \mathbf{U}_t \Sigma_t \mathbf{V}_t^\top$
 $\mathbf{D}_t = \text{diag}[\sigma_1(\sigma_1^2 + \lambda)^{-1}, \dots, \sigma_R(\sigma_R^2 + \lambda)^{-1}]$,
 $\gamma_t = \mathbf{V}_t \Sigma_t^{-1} \mathbf{D}_t \mathbf{U}_t^\top \mathbf{y}_t$
 $e_t^{(i,j)} := y_t^{(i,j)} - \langle \Phi_t^{(i,j)}, \gamma_t \rangle$
 (S2) Parallel subspace update ($r \in [R]$)
 $\mathbf{a}_r^{(1)}[t] = (1 - \mu_t \lambda / t) \mathbf{a}_r^{(1)}[t-1] + \mu_t \gamma_{t,r}^* \sum_{(i,j) \in \Omega_t} e_t^{(i,j)} (a_{j,r}^{(2)}[t-1])^* \mathbf{e}_i$
 $\mathbf{a}_r^{(2)}[t] = (1 - \mu_t \lambda / t) \mathbf{a}_r^{(2)}[t-1] + \mu_t \gamma_{t,r}^* \sum_{(i,j) \in \Omega_t} e_t^{(i,j)} (a_{i,r}^{(1)}[t-1])^* \mathbf{e}_j$
 return $(\mathbf{A}_1[t], \mathbf{A}_2[t])$
end for




Fig. 2. Multi-coil parallel MRI acquisition.

the coil sensitivity information can be further leveraged to improve the reconstruction quality. This issue is dealt with in the ensuing subsection.

B. Tomographic parallel MRI

While the approach of the previous subsection interpolates the missing k -space data and then inverts the k -space data to reconstruct the spatial domain image, one can instead pursue a tomographic approach to retrieve the images directly from partial k -space data. To generalize the measurement model (11), we adopt the multi-coil scenario where the ground-truth image \mathbf{L}_t is acquired by C coils, each sensitive to

a specific region of the image. The c -th coil's sensitivity to the image pixels is modeled by the complex-valued matrix $\mathbf{H}_c \in \mathbb{C}^{N_1 \times N_2}$. A large magnitude of $[\mathbf{H}_c]_{i,j}$ indicates a 'good view' of the (i, j) -th pixel. Ideally, the sensitivity matrices $\{\mathbf{H}_c\}_{c=1}^C$ are expected to be non-overlapping, and cover the entire image, that is $\mathbf{H}_c \odot \mathbf{H}_{c'} = \mathbf{0}$, and $\sum_{c=1}^C \mathbf{H}_c = \mathbf{1}\mathbf{1}^\top$. Let $\mathbf{L}_t^{(c)} := \mathbf{H}_c \odot \mathbf{L}_t$ denote the true image from the viewpoint of the c -th coil.

Suppose that the coil sensitivities $\{\mathbf{H}_c\}_{c=1}^C$ have been estimated directly as in e.g., [26]. With each coil c acquiring the k -space data indexed by Ω_t , our idea is to collect $C|\Omega_t|$ data per time instant t , where the (i, j) -th datum at coil c adheres to

$$y_c^{(i,j)}[t] = [\mathcal{F}(\mathbf{H}_c \odot \mathbf{L}_t)]_{i,j} + v_t^{(i,j)}, \quad (i, j) \in \Omega_t, \quad c \in [C], \quad t \in [T]. \quad (12)$$

This tomographic data model in (12) is a special case of the general one in (4) when the subsampling matrix is given by

$$\mathbf{W}_c^{(i,j)}[t] = \mathbf{H}_c \odot \mathcal{F}(\mathbf{e}_i \mathbf{e}_j^\top) \quad (13)$$

Various parallel imaging schemes have been introduced to combine the acquired images across coils for reconstructing the sought image. Among others, SENSE [26] and GRAPPA [8] are commonly used in practice; see also [6]. Each coil in the SENSE method, reconstructs an aliased image based on the subsampled k -space data (usually a fraction of phase encoding rows is selected). Then, the aliased images (or pixels) at various coils, each a linear combination of different pixels, are used to jointly reconstruct the ground-truth image [26]. Clearly, SENSE leverages the spatial diversity across coils. However, it requires knowledge of $\{\mathbf{H}_c\}_{c=1}^C$. On the other hand, the GRAPPA technique works with the raw undersampled k -space data, and interpolates the missing samples from the present neighboring ones (through a Kernel that is obtained from a calibration process using additional k -space data), and subsequently reconstructs the image [8]. The crux of GRAPPA is that the acquired k -space data per coil pertains to the ground-truth object weighted by the coil sensitivities, and thus the k -space of the object is convolved with the k -space of coil sensitivities, which smears the k -space information.

Our tomographic reconstruction approach relies on \mathbf{y}_t , which collects the complex-valued observations of coils 1 to C orderly on top of each other. Likewise, let $[\Phi_t]_{\ell,r} := [\mathcal{F}(\mathbf{H}_c \odot (\mathbf{a}_r^{(1)}[t-1] \circ \mathbf{a}_r^{(2)}[t-1]))]_{i_c, j_c}$ be a complex regression matrix comprising $C|\Omega_t|$ columns, where columns $(c' - 1)|\Omega_t| + 1$ till $c'|\Omega_t|$ correspond to coil c' . With $e_c^{(i,j)}$ denoting the fitting error for the (i, j) -th datum at coil c , due to linearity

Algorithm 3 Online sensitivity-aware and tomographic parallel MRI

input $\{y_c^{(i,j)}[t], (i,j) \in \Omega_t, c = 1, \dots, C\}_{t=1}^\infty, \{\mu_t\}_{t=1}^\infty, \lambda, R, C$.

initialize $\{\mathbf{a}_r^{(1)}[0], \mathbf{a}_r^{(2)}[0]\}_{r=1}^R$ at random.

for $t = 1, 2, \dots$ **do**

 (S1) Principal components update

$$[\phi_\ell^{(c)}]_r = \left[\mathcal{F}(\mathbf{H}_c \odot (\mathbf{a}_r^{(1)}[t-1] \circ \mathbf{a}_r^{(2)}[t-1])) \right]_{i,j}, \quad \ell \equiv (i,j)$$

$$\Phi_t = [\phi_1^{(1)}, \dots, \phi_{|\Omega_t|}^{(1)}, \dots, \phi_1^{(C)}, \dots, \phi_{|\Omega_t|}^{(C)}]^\top, \quad \Phi_t = \mathbf{U}_t \Sigma_t \mathbf{V}_t^\top$$

$$\mathbf{D}_t = \text{diag}[\sigma_1(\sigma_1^2 + \lambda)^{-1}, \dots, \sigma_R(\sigma_R^2 + \lambda)^{-1}]$$

$$\gamma_t = \mathbf{V}_t \mathbf{D}_t \mathbf{U}_t^\top \mathbf{y}_t$$

$$e_c^{(i,j)}[t] := y_c^{(i,j)}[t] - \langle \phi_\ell^{(c)}, \gamma_t \rangle$$

(S2) Parallel subspace update ($r \in [R]$)

$$\Theta[t] = \sum_{c=1}^C \mathbf{H}_c^* \odot \mathcal{F}^{-1}(\Xi_c[t])$$

$$\mathbf{a}_r^{(1)}[t] = (1 - \lambda\mu_t/t) \mathbf{a}_r^{(1)}[t-1] + \mu_t \gamma_{t,r}^* \Theta[t] (\mathbf{a}_r^{(2)}[t-1])^*$$

$$\mathbf{a}_r^{(2)}[t] = (1 - \lambda\mu_t/t) \mathbf{a}_r^{(2)}[t-1] + \mu_t \gamma_{t,r}^* \Theta^\top[t] (\mathbf{a}_r^{(1)}[t-1])^*$$

return $\{\mathbf{a}_r^{(1)}[t], \mathbf{a}_r^{(2)}[t]\}_{r=1}^R$
end for

of the Hadamard product one can simplify the gradient update using the fact that

$$\Theta[t] := \sum_{c=1}^C \sum_{(i,j) \in \Omega_t} e_c^{(i,j)}[t] (\mathbf{W}_c^{(i,j)}[t])^* = \sum_{c=1}^C \mathbf{H}_c^* \odot \mathcal{F}^{-1}(\Xi_c[t]) \quad (14)$$

where $[\Xi_c[t]]_{i,j} := e_c^{(i,j)}[t]$, $(i,j) \in \Omega_t$, and zero otherwise.

Following the general steps in Algorithm 1, the iterations of the novel parallel MRI scheme with estimated coil sensitivities, are listed under Algorithm 3.

V. SUBSPACE LEARNING VIA RANDOMIZED SUBSAMPLING

While the missing data is often attributed to outliers or malfunctioning of the acquisition process, one can purposely skip data to either facilitate the acquisition process, or lower the computational burden. This well aligned with recent efforts towards accelerating the long MRI scans, which create artifacts especially when imaging moving objects. Imagine for instance the MR scanner knowing a priori the best minimal subset of k -space data to collect per cardiac snapshot. It has then sufficient time to acquire important samples before the heart moves to a new state.

In essence, design of the optimal subsampling needs knowledge of the underlying unseen physical phenomenon of interest, that is practically infeasible. Typical sketching strategies, e.g., in the context

of matrix sparsification, assume data are *fully* available to score, and subsequently select a subset of them; see e.g., [18]. However, in MRI applications, *data acquisition* is the main challenge [14], and data are streaming as well. All in all, given the subspace estimates offered by the online iterates until time $t - 1$, our goal is to adaptively design/predict the sketching operator $\{\mathbf{W}_t^{(\ell)}\}_{\ell=1}^{L_t}$ that results in a minimal preselected sample count L_t , while attaining a prescribed reconstruction quality. In what follows, we focus on subsampling as a special sketching operator that picks only a subset of tensor entries per time instant. This subset is denoted by $\Omega_t \subset [N_1] \times \dots \times [N_M]$. The overall learning task in (P2) then amounts to tensor imputation.

A. Importance scores

Albeit streaming data poses an extra challenge to sketching (since future observations are not available), online learning offers intermediate estimates of the latent tensor subspace, namely $\{\mathbf{A}_i[t - 1]\}_{i=1}^{M-1}$, that can be leveraged to devise adaptive subsampling strategies. In order to predict Ω_t , our basic idea is to rank the samples (entries) according to their level of importance measured by a certain score along the lines of [18] and [5]. However, different from the online setup dealt with here, [18] and [5] assume batch data processing.

To gain insight, it is instructive to start with the three-way array ($M = 3$) having the entire tensor data \mathbf{Y} at hand, and seeking the factor matrices $\{\mathbf{A}_1, \mathbf{A}_2\}$. The (n_1, n_2) -th entry of t -th tensor slice can then be expressed as

$$[\mathbf{Y}_t]_{n_1, n_2} \approx \sum_{r=1}^R \gamma_{t,r} a_{n_1,r}^{(1)} a_{n_2,r}^{(2)}$$

where $\{\gamma_{t,r}\}_{r=1}^R$ are shared by every sample in the same slice t . The samples are distinguished through their weight vectors $\{a_{n_1,r}^{(1)}\}_{r=1}^R$ and $\{a_{n_2,r}^{(2)}\}_{r=1}^R$, corresponding to n_1 - and n_2 -th rows of \mathbf{A}_1 and \mathbf{A}_2 , respectively. Broadening the scope of [5], [18] to three-way arrays with PARAFAC decomposition, our proposed metric to score the (n_1, n_2) -th feature is based on the energy of the corresponding rows in the subspace matrices \mathbf{A}_1 and \mathbf{A}_2 .

According to [5] the informative samples are the ones that if missed, the matrix cannot be recovered reliably. Consider the t -th tensor slice \mathbf{Y}_t with the SVD $\mathbf{U}\Sigma\mathbf{V}^\top$. Matrix completion literature captures the information content of (i, j) -th entry through the so-termed *local coherence* measure $s_i^u := \|\mathbf{U}^\top \mathbf{e}_i\|^2$ and $s_j^v := \|\mathbf{V}^\top \mathbf{e}_j\|^2$ associated with i -th row and j -th column, respectively. In essence, the (i, j) entry is informative when both s_i^u and s_j^v are large, i.e., have large projection onto the column and row space of the matrix \mathbf{Y}_t . In particular, [5] adopts the score $s_i^u + s_j^v$ to rank entries based on their level of importance. To

gain further intuition, imagine the perturbed matrix $\mathbf{Y}_t + \mathbf{H}$ with (i, j) -th entry perturbed by $\mathbf{H} = h\mathbf{e}_i\mathbf{e}_j^\top$ for some h . A large s_i^u and s_j^v implies that \mathbf{H} is well spanned by the column and row space of \mathbf{Y}_t , and as a result \mathbf{Y}_t and $\mathbf{Y}_t + \mathbf{H}$ have the same rank. Hence, \mathbf{Y}_t can be misidentified as $\mathbf{Y}_t + \mathbf{H}$; see e.g., [5] for more technical details. Note also that in the matrix sparsification context [18] the coherence measure is referred to as the statistical leverage score which has direct ties to the residual variance of the LS estimation [18]. For a three-way array, the column vectors $\{\mathbf{a}_r^{(1)}\}_{r=1}^R$ (resp. $\{\mathbf{a}_r^{(2)}\}_{r=1}^R$) of \mathbf{A}_1 (resp. \mathbf{A}_2) span the column- (resp. row-) space of the t -th tensor slice. Albeit not necessarily orthonormal, \mathbf{A}_1 and \mathbf{A}_2 play a role similar to the orthonormal factors \mathbf{U} and \mathbf{V} .

Along this line of thought, consider now the online setup where at time instant t one has access to the subspace estimate $(\mathbf{A}_1[t-1], \mathbf{A}_2[t-1])$, and aims to acquire a few samples from the next slice \mathbf{Y}_t indexed by Ω_t . Suppose also that slices $\{\mathbf{Y}_t\}$ change slowly over time; this is the case for instance in dynamic cardiac MRI where different slices correspond to different snapshots of a patient's beating heart. Assume further that $(\mathbf{A}_1[t-1], \mathbf{A}_2[t-1])$ provide a *reliable* estimate of the underlying tensor subspace e.g., as a result of a warm initialization. It is then reasonable to adopt the metric $\|\mathbf{A}_1^\top[t-1]\mathbf{e}_{n_1}\|_F^2 + \|\mathbf{A}_2^\top[t-1]\mathbf{e}_{n_2}\|_F^2$ to predict the information content of the (n_1, n_2) -th feature at time t .

Apparently, if \mathbf{Y}_t contains innovation, not captured by the subspace $(\mathbf{A}_1[t-1], \mathbf{A}_2[t-1])$ that is learned from past data $\{\mathbf{Y}_\tau\}_{\tau=1}^{t-1}$, the informative samples may be misidentified. To cope with this issue, and avoid sampling bias due to initialization, our proposed subsampling strategy is randomized as discussed next.

B. Randomized sketching

Normalize columns of $\mathbf{A}_1[t-1]$ and $\mathbf{A}_2[t-1]$ to end up with $\bar{\mathbf{A}}_1[t-1]$ and $\bar{\mathbf{A}}_2[t-1]$, respectively. For notational brevity, let also $\bar{\mathbf{A}}_1[t-1] := [\bar{\alpha}_1^{(1)}, \dots, \bar{\alpha}_{N_1}^{(1)}]^\top$, and $\bar{\mathbf{A}}_2[t-1] := [\bar{\alpha}_1^{(2)}, \dots, \bar{\alpha}_{N_2}^{(2)}]^\top$, and score the sample $(n_1, n_2) \in [N_1] \times [N_2]$ at time t using

$$s_t(n_1, n_2) := \frac{1}{R(N_1 + N_2)} \left(\|\bar{\alpha}_{n_1}^{(1)}\|^2 + \|\bar{\alpha}_{n_2}^{(2)}\|^2 \right). \quad (15)$$

The scores $\{s_t(n_1, n_2)\}$ are positive valued and sum up to unity; thus, one can interpret them as a probability distribution over the entries. For a prescribed maximum sample count K , one can then draw K random trials from the distribution s_t to collect the important samples in the set Ω_t .

Sampling can be performed with, or, without replacement depending on how the image energy is distributed. Sampling without replacement results in exactly K samples. However, sampling with replacement can even reduce the sample count ($|\Omega_t| < K$) when the samples are nonuniformly informative. For MRI images typically most of the energy is concentrated in low-frequency k -space samples. Thus, it

Algorithm 4 Random subsampling for M -way arrays

input $\{\mathbf{A}_m[t-1]\}_{m=1}^{M-1}$, K , and R

$$\bar{\mathbf{A}}_m[t-1] := \mathbf{A}_m[t-1] \text{diag}(\|\mathbf{a}_1^{(m)}[t-1]\|, \dots, \|\mathbf{a}_R^{(m)}[t-1]\|)^{-1}$$

$$\bar{\mathbf{A}}_m[t-1] := [\bar{\boldsymbol{\alpha}}_1^{(m)}, \dots, \bar{\boldsymbol{\alpha}}_{N_m}^{(m)}]^\top$$

$$s_t(n_1, \dots, n_{M-1}) := \frac{1}{R \sum_{m=1}^{M-1} N_m} \sum_{m=1}^{M-1} \|\bar{\boldsymbol{\alpha}}_{n_m}^{(m)}\|^2$$

Draw K random trials from $[N_1] \times \dots \times [N_M]$ based on s_t to form Ω_t

output Ω_t

seems more natural to consider random draws with replacement which is more likely to discard the less informative samples (with tiny score) that have only a negligible contribution to the image. It is also worth commenting that nowadays MRI scanners can accommodate quickly changing the gradient pulse sequence to acquire the sampled phase encoding lines in real time [15].

Adopting randomized sketching along with the online iterates for subspace learning in (10), the resulting procedure for a general M -way data array is listed under Algorithm 5. The iterative scheme begins with a warm initialization, obtained for instance after first running the algorithm over a small training dataset. Each iteration (time instant) t comprises three major steps, where the first step (S0) *probabilistically* decides on the subsampling set Ω_t , that is subsequently used to acquire the corresponding samples in \mathbf{Y}_t . Based on the partial samples in Ω_t , the second step (S1) finds the principal components of t -th frame across the subspace bases $\{\mathbf{a}_r^{(1)}[t-1] \mathbf{a}_r^{(2)\top}[t-1]\}_{r=1}^R$. The innovation of the new (imputed) datum captured through the error term $\{e_t^{(i,j)}\}_{(i,j) \in \Omega_t}$, in the third step (S2) then refines the subspace bases.

One important question at this point pertains to the (average) number of samples acquired per slice by the random subsampling with replacement. This depends on the degree of sample nonuniformity, and the initialization. In an extreme case with equally important entries, exactly K samples are acquired per time slot, which can be considerably lower for nonuniform entries. To see this, introduce a random variable $X_{m,n}$ denoting the frequency of choosing the (n_1, n_2) -th entry after K trials. The random sample count is then $|\Omega_t| = \sum_{n_1, n_2} \mathbb{1}_{\{X_{n_1, n_2}\}}$, where the indicator $\mathbb{1}_{\{x\}}$ takes value of one if $x > 0$, and zero otherwise. In general, the random variables $\mathbb{1}_{\{X_{n_1, n_2}\}}$ are dependent, which renders the distribution analysis for $|\Omega_t|$ formidable. The expected sample count per slice however can be expressed as

$$\mathbb{E}[|\Omega_t|] = \sum_{m=1}^M \sum_{n=1}^N \left(1 - [1 - s_t(n_1, n_2)]^K\right).$$

Finally, the average sample count across time and entries is given by $\bar{N}_t := (1/t) \sum_{\tau=1}^t \mathbb{E}[|\Omega_\tau|]$.

Algorithm 5 Randomized tensor subspace learning for imputation of M -way arrays**input** $\{\mu_t\}_{t=1}^\infty, K, R, \lambda$ **initialize** $\{\mathbf{A}_m[0]\}_{m=1}^{M-1}$ with a warm startup**for** $t = 1, \dots$ **do****(S0) Random subsampling**Acquire $\{y_t^{(n_1, \dots, n_{M-1})}\}_{(n_1, \dots, n_{M-1}) \in \Omega_t}$ based on Algorithm 1**(S1) Principal components update**

$$[\Phi_t]_{(n_1, \dots, n_{M-1}), r} = \prod_{m=1}^{M-1} a_{n_m, r}[t-1], \quad \Phi_t = \mathbf{U}_t \Sigma_t \mathbf{V}_t^\top$$

$$\mathbf{D}_t = \text{diag}[\sigma_1(\sigma_1^2 + \lambda)^{-1}, \dots, \sigma_R(\sigma_R^2 + \lambda)^{-1}],$$

$$\gamma_t = \mathbf{V}_t \Sigma_t^{-1} \mathbf{D}_t \mathbf{U}_t^\top \mathbf{y}_t$$

$$e_t^{(n_1, \dots, n_{M-1})} := y_t^{(n_1, \dots, n_{M-1})} - \langle \phi_t^{(n_1, \dots, n_{M-1})}, \gamma_t \rangle$$

(S2) Parallel subspace update $[(m, r) \in [M] \times [R]]$

$$\mathbf{a}_r^{(m)}[t] = (1 - \mu_t \lambda / t) \mathbf{a}_r^{(m)}[t-1]$$

$$+ \mu_t \gamma_{t,r}^* \sum_{(n_1, \dots, n_M) \in \Omega_t} e_t^{(n_1, \dots, n_{M-1})} \prod_{m=1, m \neq m}^{M-1} (a_{n_m, r}^{(m)}[t-1])^* \mathbf{e}_m$$

end for**output** $\{\mathbf{A}_1[t], \mathbf{A}_2[t]\}$

VI. NUMERICAL TESTS

Performance of the novel tensor subspace learning scheme for online MRI reconstruction is assessed in this section using real cardiac MRI data. Two datasets are adopted as follows: (D1) single-coil acquisition of 256 cardiac cine MRI frames, each of size 200×256 pixels; and, (D2) 16-coil dataset including 26 frames of size 192×120 acquired by the Center for Magnetic Resonance Research (CMRR) at the University of Minnesota. The experiments were run using Matlab R2016a on a Ubuntu 15.10 machine with 64GB RAM. The processors include an Intel core i7 CPU, and an NVIDIA GDX 970 GPU.

For validation purposes the k -space data is fully acquired. To simulate the real-world undersampled data, as it is common with practical MRI scanners, variable density Cartesian sampling is used to randomly sample a small fraction π of the phase encoding lines. They are sampled based on the polynomial distribution $p_i = i^\alpha / \sum_{i=1}^{M/2-1} i^\alpha$, where i is the distance from the center line. The center line ($i = 0$) is always chosen, and we fix $\alpha = -1$, so that the low frequency components carrying most of the image energy are more likely to be selected.

A. Real-time reconstruction

The frequency-domain interpolation-based scheme in Algorithm 2 is tested on the dataset (D1) forming the tensor $\underline{\mathbf{X}} \in \mathbb{C}^{200 \times 256 \times 256}$. The singular values of the unfolded tensor along the temporal axis, as well as the k_x and k_y axes are plotted in Fig. 3. The unfolded k -space data matrix along the temporal dimension exhibits only about four dominant singular values, indicating high spatio-temporal correlations indicative of low rank for tensor $\underline{\mathbf{X}}$. The first 5 scans are fully acquired, and used as training data to provide a warm startup for the tensor subspace initialization. In practice, such a warm initialization can be even provided by historical data of other patients, or, a complementary prescan of the same patient. For the rest of 251 frames, partial k -space data are acquired based on the variable density Cartesian sampling.

Upon choosing the step size $\mu_t = 0.01$ and $\lambda = 2$, we plot time evolution of the normalized mean-square error $\text{NMSE}_t = \|\mathbf{X}_t - \hat{\mathbf{X}}_t\|_F^2 / \|\mathbf{X}_t\|_F^2$ with \mathbf{X}_t and $\hat{\mathbf{X}}_t$ denoting, respectively, the true and estimated frames per t -th frame. As a subjective metric, structural similarity index (SSIM) [31] is also plotted over time in Fig. 4 under different tensor ranks and acceleration rates. The resulting reconstructed images at $t = 10, 100, 170, 240$ are also shown and compared against the gold standard (full k -space data) in Fig. 5. Apparently, the first 50 frames show a sharp improvement in learning the subspace and consequently the reconstruction accuracy, which improves gradually for the subsequent frames. More epochs over the data (compared to the single visit for real-time scheme) are needed to fully learn the subspace. It turns out that 5 epochs suffice to learn the subspace, where the output of each epoch serves to initialize the subspace for the next one. The reconstruction error for this offline scheme is also shown in Fig. 4 as the benchmark. Note that under 10-fold acceleration, the acquisition time for each frame is about xxxx sec while for $R = 100$ the reconstruction takes about 0.1 sec on a GPU. This can be further shorten to 0.05 sec with $R = 50$ that yields relatively close reconstruction quality; NMSE=0.029, 0.03 for $R = 100, 50$, respectively. Therefore, the novel scheme can reconstruct 20 frames per second, which in turn makes *real-time* reconstruction practically feasible.

B. Comparisons

Online MRI approaches so far include the Kalman filtering based ones [20], [28], and the more recent compressive sampling (CS) based versions e.g., in [21], [30], [33]. The Kalman filtering method enjoys relatively fast reconstruction at the expense of low image quality that CS schemes improve upon, at the expense of slower reconstruction that is due to the non-smooth regularization. Indeed, existing




Fig. 3. Singular values of the matrix each corresponding to the unfolded k -space tensor across a different dimension.




Fig. 4. Evolution of NMSE (left) and SSIM (right) over iteration index (time) when a fraction π of phase encoding lines are sampled ($R = 100$, $\lambda = 2$, $\mu = 0.01$).

online reconstruction schemes are not real time since the reconstruction time is typically slower than the acquisition time.

To evaluate the merits of our novel reconstruction scheme, we compare with the differential CS scheme [21] as an state-of-the-art alternative. The crux of [21] is that the difference between subsequent




Fig. 5. Real-time reconstruction of *in vivo* MRI dataset based on Algorithm 2 at different time points $t = 10, 100, 170, 240$ (left to right) under (top) 10-fold, and (bottom) 4-fold acceleration, when $R = 100$. A patch of size 138×100 pixels is shown.

frames is sparse. Let \mathbf{X}_{t-1} be an estimate of the previous frame $t - 1$. To reconstruct \mathbf{X}_t the difference frame $\tilde{\mathbf{X}} := \mathbf{X}_t - \mathbf{X}_{t-1}$ is assumed sparse, and can thus be recovered by solving the non-smooth LASSO [29] program

$$\tilde{\mathbf{X}}_t = \arg \min_{\tilde{\mathbf{X}}} \sum_{(i,j) \in \Omega_t} (y_t^{(i,j)} - [\mathcal{F}(\mathbf{X}_{t-1})]_{i,j} - [\mathcal{F}(\tilde{\mathbf{X}})]_{i,j})^2 + \lambda \|\tilde{\mathbf{X}}\|_1. \quad (16)$$

Regarding the differential CS, it is apparently slow as it demands solving a LASSO program per time instant. Additionally, the error can gradually accumulate over time. SpaRSA package [32] is used to return the LASSO solution in (16) where for the stop criterion we choose the duality gap to be less than 0.01. Each iteration is initialized with the difference image obtained in the previous iteration as a warm start up, and the maximum number of iterations is confined to 100.

Temporal profile of the reconstructed images along the oblique lines are shown in Fig. 6 for 4-fold and 10-fold accelerations with variable density Cartesian sampling. Parameter choices were $\lambda = 0.001$ for differential CS, and $\mu = 0.01$, $\lambda = 2$, and $R = 100$ for our novel method. As pinpointed by the arrows, differential CS leads to temporal blurring artifacts, while the proposed scheme can track the changes




Fig. 6. Temporal profile of the oblique lines on the left side for the gold standard, differential CS, and the proposed TSL method for 4-fold (top) and 10-fold (bottom) acceleration under variable density sampling when $R = 100$ and only the first five frames are used for training. The pixels on the lines are arranged vertically in the profile images.

more accurately, and thus reveals the detailed temporal changes. Table I also lists the lower reconstruction time for TSL scheme relative to the differential CS scheme. This becomes more pronounced for higher acceleration rates as the operation count $\mathcal{O}(|\Omega_t|R^2)$ decreases (c.f. Remark 1), while the degree of non-smoothness for the LASSO program (16), and consequently the convergence time, increases. Note that more than 90% of the TSL runtime per iteration stems from solving (exactly) the ridge-regression task (S1) to update γ_t , which can be further reduced if one resorts to an inexact solution.

C. Random k -space subsampling

The random subsampling policy developed in Section V is tested here to sample Cartesian phase encoded lines for interpolation of missing k -space data in dataset (D1). Since only the rows are se-

TABLE I
AVERAGE FRAME RECONSTRUCTION TIME (SECONDS) FOR TSL AND DIFFERENTIAL CS SCHEMES.

Acceleration	Diff. CS	TSL, $R = 100$ (CPU)	TSL, $R = 100$ (GPU)	TSL, $R = 50$ (GPU)
10-fold	0.55	0.14	0.11	0.05
4-fold	0.56	0.28	0.23	0.1

lected, the score in (15) is marginalized over the columns to arrive at the modified score $s_t(n_1) = \frac{1}{R(N_1+N_2)}(N_2\|\bar{\alpha}_{n_1}^{(1)}\|^2 + R)$, and subsequently random trials are drawn with replacement to sample the rows. As before, the first five frames are fully acquired to serve as a warm subspace startup. Evolution of NMSE is depicted in Fig. 7 under 10-fold acceleration for the adaptive sampling scheme, and its non-adaptive counterpart using variable density sampling with parameters $\alpha = -1, -0.5$ for the polynomial distribution p_i . After dozen iterations the latent structure in k -space data is gradually learned and as a result subspace-driven sampling starts outperforming the variable density sampling. This observation suggests that one better adopt variable density sampling with $\alpha = -1$ to pick the low frequency components for the early iterations to end up with a reliable subspace estimation, and then switch to adaptive subsampling for improved quality. The reconstructed frames at time $t = 100$ are shown in Fig. 9, where one can confirm from the residual images that our novel scheme creates less artifacts as delineated by the arrows. For the resultant sampling patterns with 256 temporal realizations, the sampling distribution for phase encoded lines is also depicted in Fig. 8.

D. Batch MRI processing

The recursive structure and the lightweight iterations offered by the TSL scheme suit batch processing of large MRI datasets. Inspired by the incremental gradient methods, one can allow multiple epochs over the data, where the subspace learnt in the first epoch initializes the second epoch (with possibly random re-ordering of the frames), and so on. Batch performance of TSL is compared against the k - t FOCUSS scheme [10] that achieves the reconstruction quality of k - t SPARSE [17], while being computationally more appealing as it relies on a successive quadratic optimization. For TSL, we choose $R = 75$ and allow 4 epochs, while for the k - t FOCUSS we use the publicly available source code provided in [10] with 2 outer iterations and 40 inner iterations. Runtime (seconds) is compared for the k - t FOCUSS and TSL schemes in Table II, where it is evident that the TSL scheme converges much faster especially when running on GPU. Note that k - t FOCUSS cannot be implemented on GPU for the provided code




Fig. 7. NMSE of the proposed score-based adaptive phase encoding versus the variable density counterpart with 10-fold acceleration and $R = 100$.




Fig. 8. Sampling distribution ρ for data driven (adaptive) and variable density schemes when $\pi = 0.1$.

since each iteration involves the large complex-valued tensor $\underline{\mathbf{X}} \in \mathbb{C}^{200 \times 256 \times 256}$, which exceeds the GPU memory.




Fig. 9. Real-time reconstruction comparison at time instant $t = 100$ for differential CS and TSL schemes under variable density and adaptive sampling. Top (bottom) row corresponds to 10-fold (4-fold) acceleration. first column: original frame; second column: differential CS; third column: TSL with variable density sampling; and fourth column: TSL with adaptive sampling.

TABLE II
RUNTIME (SEC) AND AVERAGE NMSE FOR BATCH PROCESSING OF k - t TSL AND FOCUSS [10] UNDER 10-FOLD ACCELERATION.

Scheme	Avg. NMSE	Runtime (CPU)	Runtime (GPU)
FOCUSS	0.049	340	N/A
TSL	0.01	84	27.75

E. Tomographic parallel MRI

The multi-coil dataset (D2) is used to evaluate performance for parallel MRI reconstruction. The coil sensitivity maps are estimated using the sum-of-squares method from the first k -space frame that is




Fig. 10. Coil sensitivity magnitude (top) and phase (bottom) for eight coils in parallel MRI dataset D2 with $R = 200$ and $\lambda = 1$. The phase is normalized by π .

TABLE III

NMSE OF THE TOMOGRAPHIC PARALLEL MRI UNDER VARIOUS ACCELERATION FACTORS WHEN $R = 200$.

Acceleration	1	2	4	10
NMSE	0.0029	0.0031	0.0038	0.0067

acquired fully. Phase smoothing via polynomial fitting as in [26] is then used to smooth the sharp phase transitions. The resulting phase and magnitude for the first 8 coils are depicted in Fig. 10. Tensor rank was chosen to be $R = 200$. Due to limited number of temporal slices ($T = 26$), the batch procedure is adopted with multiple visits over data. Selecting $\lambda = 1$, two representative reconstructed images under various acceleration factors are depicted in Fig. 11. Apparently, the reconstructed images qualitatively look quite close to the ground-truth one, where for rate 10 acceleration some flow artifacts (as pinpointed by arrows) begin appearing around the heart. This is also quantitatively confirmed by the small average NMSE listed in Table III.

VII. CONCLUSIONS AND FUTURE DIRECTIONS

Fast and effective analytics are developed in this paper for real-time reconstruction of MR images. Treating a temporal image sequence as a multi-way data array, a novel tensor subspace learning approach was introduced based on the PARAFAC decomposition. The correlation structure across various dimensions is learned from the tensor's latent subspace. A rank-regularized least-squares estimator was put forth to learn this tensor subspace when a Tykhonov-type regularizer which promotes low rank for




Fig. 11. Reconstructed parallel MRI images under different acceleration factors for two representative images corresponding to top and bottom rows when $R = 200$ and $\lambda = 1$.

the data tensor. Adopting stochastic alternating minimization, recursive algorithms were developed to track the subspace and reconstruct the image 'on the fly.' The resulting algorithm enjoys lightweight iterations with parallelized computations, which makes it attractive for high-resolution real-time MRI. Leveraging the online subspace iterates offered by the algorithm, adaptive subsampling strategies were also devised to randomly predict informative samples for acquisition of subsequent k -space frames. GPU-based simulated tests with real cardiac cine MRI datasets corroborated the effectiveness of the novel approach. To broaden the scope of the present work there are several intriguing questions to address including the extension and evaluation of the novel analytics for volumetric and 4D MRI with more sophisticated such as radial sampling trajectories. Another important direction pertains to performance analysis of the adaptive subsampling policy for selecting the best set of k -space samples.

VIII. ACKNOWLEDGMENTS

We would like to acknowledge Dr. Sebastian Schmitter from the Center for Magnetic Resonance Research, Minneapolis, Minnesota, for providing us with the dataset (D2) and helpful discussions. The valuable comments by Dr. Leslie Ying from the State University of New York at Buffalo who provided us with the dataset (D2) are also gratefully appreciated.

REFERENCES

- [1] E. Acar, D. M. Dunlavy, T. G. Kolda, and M. Mrup, "Scalable tensor factorizations for incomplete data," *Chemometrics and Intelligent Laboratory Systems*, vol. 106, no. 1, pp. 41–56, Mar. 2011.
- [2] B. Baingana and G. B. Giannakis, "Joint community and anomaly tracking in dynamic networks," *IEEE Trans. on Signal Processing*, vol. 68, no. 8, pp. 2013–2025, Apr. 2016.
- [3] J. A. Bazerque, G. Mateos, and G. B. Giannakis, "Rank regularization and Bayesian inference for tensor completion and extrapolation," *IEEE Trans. on Signal Process.*, vol. 61, no. 22, pp. 5689–5703, Nov. 2013.
- [4] J. M. F. T. Berge and N. D. Sidiropoulos, "On uniqueness in CANDECOMP/PARAFAC," *Psychometrika*, vol. 67, no. 3, pp. 399–409, Sep. 2002.
- [5] Y. Chen, S. Bhojanapalli, S. Sanghavi, and R. Ward, "Coherent matrix completion," in *Proc. of 31th International Conference on Machine Learning*, Beijing, China, Jun. 2014, pp. 674–682.
- [6] A. Deshmane, V. Gulani, M. A. Griswold, and N. Seiberlich, "Parallel MR Imaging," *Journal of Magnetic Resonance Imaging*, vol. 36, no. 1, pp. 55–72, Jul. 2012.
- [7] M. Fazel, "Matrix rank minimization with applications," Ph.D. dissertation, Stanford University, 2002.
- [8] M. A. Griswold, P. M. Jakob, R. M. Heidemann, M. Nittka, V. Jellus, J. Wang, B. Kiefer, and A. Haase, "Generalized autocalibrating partially parallel acquisitions (GRAPPA)," *Magnetic Resonance in Medicine*, vol. 47, no. 6, pp. 1202–1210, Jun. 2002.
- [9] F. Huang, J. Akao, S. Vijayajumar, G. R. Duensing, and M. Limkeman, "k-t GRAPPA: A k-space implementation for dynamic MRI with high reduction factor," *Magnetic Resonance in Medicine*, vol. 54, no. 5, pp. 1172–1184, Nov. 2005.
- [10] H. Jung, K. Sung, K. S. Nayak, E. Y. Kim, and J. C. Ye, "k-t FOCUSS: A general compressed sensing framework for high resolution dynamic MRI," *Magnetic Resonance in Medicine*, vol. 61, no. 1, pp. 103–116, Jan. 2009.
- [11] T. G. Kolda and B. W. Bader, "Tensor decompositions and applications," *SIAM Review*, vol. 51, no. 3, pp. 455–500, Aug. 2009.
- [12] S. Kozerke, J. Tsao, R. Razavi, and P. Boesiger, "Accelerating cardiac cine 3D imaging using k-t BLAST," *Magnetic Resonance in Medicine*, vol. 52, no. 1, pp. 19–26, Jul. 2004.
- [13] J. Kruskal, "Three-way arrays: Rank and uniqueness of trilinear decompositions with application to arithmetic complexity and statistics," *Lin. Alg. Applications*, vol. 18, no. 2, pp. 95–138, 1977.
- [14] P. C. Lauterbur, "Image formation by induced local interactions: Examples employing nuclear magnetic resonance," *Nature*, vol. 242, no. 5394, pp. 190–191, Mar. 1973.
- [15] Z.-P. . Liang and P. C. Lauterbur, *Principles of Magnetic Resonance Imaging: A Signal Processing Perspective*. The Institute of Electrical and Electronics Engineers Press, 2000.

- [16] S. G. Lingala, Y. Hu, E. DiBella, and M. Jacob, "Accelerated dynamic MRI exploiting sparsity and low-rank structure: k-t SLR," *IEEE Trans. on Medical Imaging*, vol. 30, no. 5, pp. 1042–1054, May 2011.
- [17] M. Lustig, J. M. Santos, D. L. Donoho, and J. M. Pauly, "k-t SPARSE: High frame rate dynamic MRI exploiting spatio-temporal sparsity," in *Proc. of the 13th ISMRM Annual Meeting and Exhibition*, vol. 2420, Seattle, May 2006.
- [18] M. W. Mahoney, "Randomized algorithms for matrices and data," *Foundations and Trends® in Machine Learning*, vol. 3, no. 2, pp. 123–224, 2011.
- [19] J. Mairal, F. Bach, J. Ponce, and G. Sapiro, "Online learning for matrix factorization and sparse coding," *J. of Machine Learning Research*, vol. 11, pp. 19–60, Jan. 2010.
- [20] A. Majumdar, N. Ansari, H. Aggarwal, and P. Biyani, "Impulse denoising for hyper-spectral images: A blind compressed sensing approach," *Elsevier Signal Processing*, vol. 119, pp. 136–141, Feb. 2016.
- [21] A. Majumdar, R. K. Ward, and T. Aboulnasr, "Compressed sensing based real-time dynamic MRI reconstruction," *Medical Imaging, IEEE Trans. on*, vol. 31, no. 12, pp. 2253–2266, Dec. 2012.
- [22] M. Mardani, G. Mateos, and G. B. Giannakis, "Dynamic anomalography: Tracking network anomalies via sparsity and low rank," *IEEE Journal of Selected Topics on Signal Processing*, vol. 7, no. 1, pp. 50–66, Feb. 2013.
- [23] —, "Subspace learning and imputation for streaming big data matrices and tensors," *IEEE Trans. on Signal Processing*, vol. 63, no. 10, pp. 2663–2677, May 2015.
- [24] R. Otazo, L. Feng, H. Chandarana, T. Block, L. Axel, and D. K. Sodickson, "Combination of compressed sensing and parallel imaging for highly-accelerated dynamic MRI," in *Proc. 9th IEEE International Symposium on Biomedical Imaging*, Barcelona, Spain, May 2012, pp. 980–983.
- [25] H. Pedersen, S. Kozerke, S. Ringgaard, K. Nehrke, and W. Y. Kim, "k-t PCA: Temporally constrained k-t BLAST reconstruction using principal component analysis," *Magnetic Resonance in Medicine*, vol. 62, no. 3, pp. 706–716, Sep. 2009.
- [26] K. P. Pruessmann, M. Weiger, M. B. Scheidegger, P. Boesiger *et al.*, "SENSE: sensitivity encoding for fast MRI," *Magnetic Resonance in medicine*, vol. 42, no. 5, pp. 952–962, Nov. 1999.
- [27] S. Ravishanker and Y. Bresler, "MR image reconstruction from highly undersampled k-space data by dictionary learning," *IEEE Trans. on Medical Imaging*, vol. 30, no. 5, pp. 1028–1041, May 2011.
- [28] U. Sumbul, J. Santos, and J. Pauly, "A practical acceleration algorithm for real-time imaging," *IEEE Trans. on Medical Imaging*, vol. 12, no. 28, pp. 2042–2051, 2009.
- [29] R. Tibshirani, "Regression shrinkage and selection via the lasso," *Journal of the Royal Statistical Society. Series B (Methodological)*, pp. 267–288, 1996.
- [30] N. Vaswani, "LS-CS-residual (LS-CS): compressive sensing on least squares residual," *IEEE Trans. on Signal Processing*, vol. 58, no. 8, pp. 4108–4120, Aug. 2010.
- [31] Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli, "Image quality assessment: from error visibility to structural similarity," *IEEE Trans. on Image Processing*, vol. 13, no. 4, pp. 600–612, Apr. 2004.
- [32] S. J. Wright, R. D. Nowak, and M. A. Figueiredo, "Sparse reconstruction by separable approximation," *IEEE Trans. on Signal Processing*, vol. 57, no. 7, pp. 2479–2493, 2009.
- [33] Y. Yu, J. Jin, F. Liu, and S. Crozier, "Multidimensional compressed sensing mri using tensor decomposition-based sparsifying transform," *PloS One*, vol. 9, no. 6, p. e98441, Jun. 2014.